



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
-----------------	-------------	----------------------	---------------------	------------------

10/750,859

01/05/2004

Kenichi Miki

501.43387X00

3962

24956

7590

10/03/2005

MATTINGLY, STANGER, MALUR & BRUNDIDGE, P.C.
1800 DIAGONAL ROAD
SUITE 370
ALEXANDRIA, VA 22314

EXAMINER

HICKS, MICHAEL J

ART UNIT

PAPER NUMBER

2165

DATE MAILED: 10/03/2005

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

10/750,859

Applicant(s)

MIKI ET AL.

Examiner

Michael J. Hicks

Art Unit

2165

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 24 February 2004.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-6 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-6 is/are rejected.
- 7) ☒ ~~Claim(s) 2 is/are objected to.~~
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 05 January 2004 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
- Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
- Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☒ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☒ Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
- Paper No(s)/Mail Date 01/05/04; 06/29/05

- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____
- 5) ☐ Notice of Informal Patent Application (PTO-152)
- 6) ☐ Other: _____


SAM RIMELL
PRIMARY EXAMINER

AT

DETAILED ACTION

Claim Objections

1. Claim 2 objected to because of the following informalities: In line two of Claim 2, the phrase "in which the reception tines" should read "in which the reception times".

Appropriate correction is required.

2. Claims 1 and 5 objected to because of the following informalities: In reference to the term "superior device", there is no mention as to what the device is meant to be superior to, or in what way it is superior. Examiner therefor finds the term vague and in need of clarification.

Appropriate correction is required.

Claim Rejections - 35 USC § 102

3. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

4. Claims 1 and 4 rejected under 35 U.S.C. 102(e) as being anticipated by Brown (U.S. Pre Grant Publication Number 2003/0014598).

As per Claim 1, Brown discloses a storage device system in a computer system having a plurality of superior devices and a plurality of storage device systems for receiving write-in data from at least one of the superior devices, comprising: a physical storage device, responsive to a logical volume which is controlled so that identical data is saved across the plurality of storage device systems, for storing data on the logical volume (i.e. *"Volume managers are primarily used to organize storage devices into logical volumes, which may span multiple storage devices...Briefly, the present invention includes at least a first and a second server of a cluster of servers being equipped with complementary software RAID drivers and distributed lock managers that enable the first server to delegate to the second server, writing of a version of a unit of coherent data into a number of storage devices coupled to the server cluster."* The preceding text excerpt clearly indicates that a cluster of servers/a plurality of superior devices and a second/delegated server and number of storage devices/a plurality of storage device systems exist that include a storage device/physical storage device which is responsive to a logical volume which is controlled with complementary software RAID drivers/controlled so that identical data is saved across the plurality of storage devices. Note that RAID arrays are used to save identical data across a plurality of storage devices.) (Page 1, Paragraphs 4, 15); a device for saving the time of reception on which write-in data was received from a superior device (i.e. *"...in performing a delegated write, the delegated server may obtain at least a shared read lock on the unit of coherent data and validating a timestamp of the version of the unit of coherent data to be written."* The preceding text excerpt clearly indicates that delegated server obtains/validates/saves a timestamp/the time of reception on which write-in data was received from a first/ server/delegating server/superior device) (Pg 1, Paragraph 18); a communication interface device for transmitting write-in data addressed to the logical volume and a corresponding reception time to another storage device system and for receiving write-in data and corresponding reception time from the

storage device system (i.e. *"...in performing a delegated write, the delegated server may obtain at least a shared read lock on the unit of coherent data and validating a timestamp of the version of the unit of coherent data to be written...the LastWriteTimeStamp property in the ManagedObject is the time at which the object was last written to the storage devices... FIGS. 4c and 4d show the delayed write to storage devices."*) The preceding text excerpt clearly indicates that the delegated server validates the timestamp/reception time of the unit of data to be written. In order to validate the timestamp the delegating server must first receive it. Also note that because any of the cluster of servers can be designated as a delegating, or delegated server the server has the capacity to both transmit and receive the timestamp information. This property is demonstrated pertaining to the storage devices/storage device systems through the reference to the LastUpdateWriteStamp. The text excerpt also indicates, along with figures 4c and 4d, that write-in data may also be transferred in this way.) (Figures 4c, 4d; Pg 1, Paragraph 18; Page 3, Paragraph 46; Page 4, Paragraph 69); and a data consistency holding control device for effecting control to write write-in data which was received from the superior device and write-in data which was received through the communication interface device into the physical storage device after such write-in data has been made to wait in a temporary storage device for more than predetermined time from the reception time corresponding to the write-in data to the logical volume (i.e. *"FIGS. 4c and 4d show the delayed write to storage devices. This is typically referred to as a write-back operation. FIG. 4d shows the writeback as done by the node with the working instance of the data blocks..."*) The preceding text excerpt clearly indicates that the data blocks/write in data is written to storage devices after a delay/a predetermined time from the reception time spent in Node Y/temporary storage.) (Figure 4c; Figure 4d; Page 4, Paragraph 69), so that, when write-in data which was received from the superior device and write-in data which was received through the communication interface device are written in an overlapped manner into the same storage location of the physical storage device (note that such overlapped writing is also a feature of RAID), they

are written in the order of the reception time thereof (i.e. "...*timestamp aging, as well as other techniques may be employed.*" The preceding text excerpt clearly indicates that timestamp aging, which assigns priority on a first come first serve basis, is used (e.g. they are written in the order of reception).) (Page 4, Paragraph 70).

As per Claim 4, Brown discloses the storage device system further has a device for receiving a request for locking a partial region of the logical volume from the superior device and for locking the partial region, a device for transmitting the locking request which was received through the communication interface device to the other storage device system, a device for receiving the locking request through the communication interface device from the other storage device system and for locking a designated partial region and a device for rejecting a request for write-in of data from the superior device and the other storage device system to the partial region, except for a case in which it is a request from the superior device in which the partial region was locked (i.e. "*Once the writeback operation is started, the software RAID driver secures an exclusive write lock on the stripe.*" The preceding text excerpt clearly indicates that in order to secure the exclusive lock for the stripe/partial region, the second server/delegated server/storage system device must utilize the software RAID driver to request an exclusive lock from the first server/delegating server/superior device and then receive that request back in order to transmit the request to the logical volume, therefor locking the stripe/partial region. Note that an exclusive lock would deny access to the stripe to any system except the system which currently holds the lock.) (Page 4, Paragraph 71).

Claim Rejections - 35 USC § 103

5. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

Art Unit: 2165

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

6. Claims 2, and 5-6 rejected under 35 U.S.C. 103(a) as being unpatentable over Brown (U.S. Pre Grant Publication Number 2003/0014598) in view of Azagury et al. ("Advanced functions for storage subsystems: Supporting continuous availability", IBM Systems Journal, Volume 42, Number 2; Internet Publication Date 5/1/2003, and referred to hereinafter as Azagury).

As per Claim 2, Brown discloses a storage device system in a computer system having a plurality of superior devices and a plurality of storage device systems for receiving write-in data from at least one of the superior devices, comprising: a physical storage device, responsive to a logical volume which is controlled so that identical data is saved across the plurality of storage device systems, for storing data on the logical volume (i.e. *"Volume managers are primarily used to organize storage devices into logical volumes, which may span multiple storage devices....Briefly, the present invention includes at least a first and a second server of a cluster of servers being equipped with complementary software RAID drivers and distributed lock managers that enable the first server to delegate to the second server, writing of a version of a unit of coherent data into a number of storage devices coupled to the server cluster."* The preceding text excerpt clearly indicates that a cluster of servers/a plurality of superior devices and a second/delegated server and number of storage devices/a plurality of storage device systems exist that include a storage device/physical storage device which is responsive to a logical volume which is controlled with complementary software RAID drivers/controlled so that identical data is saved across the plurality of storage devices. Note that RAID arrays are used to save identical data across a plurality of storage devices.) (Page 1, Paragraphs 4, 15); a device for saving the time of reception on which write-in data was received from a superior device (i.e. *"...in performing a delegated write, the delegated server may obtain at least a shared read lock on the unit of coherent data and*

Art Unit: 2165

validating a timestamp of the version of the unit of coherent data to be written." The preceding text excerpt clearly indicates that delegated server obtains/validates/saves a timestamp/the time of reception on which write-in data was received from a first/ server/delegating server/superior device) (Pg 1, Paragraph 18); a communication interface device for transmitting write-in data addressed to the logical volume and a corresponding reception time to another storage device system and for receiving write-in data and corresponding reception time from the storage device system (i.e. *"...in performing a delegated write, the delegated server may obtain at least a shared read lock on the unit of coherent data and validating a timestamp of the version of the unit of coherent data to be written...the LastWriteTimeStamp property in the ManagedObject is the time at which the object was last written to the storage devices... FIGS. 4c and 4d show the delayed write to storage devices."* The preceding text excerpt clearly indicates that the delegated server validates the timestamp/reception time of the unit of data to be written. In order to validate the timestamp the delegating server must first receive it. Also note that because any of the cluster of servers can be designated as a delegating, or delegated server the server has the capacity to both transmit and receive the timestamp information. This property is demonstrated pertaining to the storage devices/storage device systems through the reference to the LastUpdateWriteStamp. The text excerpt also indicates, along with figures 4c and 4d, that write-in data may also be transferred in this way.) (Figures 4c, 4d; Pg 1, Paragraph 18; Page 3, Paragraph 46; Page 4, Paragraph 69); and a data consistency holding control device for effecting control to write write-in data which was received from the superior device and write-in data which was received through the communication interface device into the physical storage device after such write-in data has been made to wait in a temporary storage device for more than predetermined time from the reception time corresponding to the write-in data to the logical volume (i.e. *"FIGS. 4c and 4d show the delayed write to storage devices. This is typically referred to as a write-back operation. FIG. 4d shows the writeback as done by the node with the working instance of the data blocks..."* The

preceding text excerpt clearly indicates that the data blocks/write in data is written to storage devices after a delay/a predetermined time from the reception time spent in Node Y/temporary storage.) (Figure 4c; Figure 4d; Page 4, Paragraph 69), so that, when write-in data which was received from the superior device and write-in data which was received through the communication interface device are written in an overlapped manner into the same storage location of the physical storage device (note that such overlapped writing is also a feature of RAID), they are written in the order of the reception time thereof (i.e. *"...timestamp aging, as well as other techniques may be employed."* The preceding text excerpt clearly indicates that timestamp aging, which assigns priority on a first come first serve basis, is used (e.g. they are written in the order of reception).) (Page 4, Paragraph 70) and a device for searching write-in data to determine for which data said more than predetermined time has passed from the reception time (i.e. *"FIGS. 4c and 4d show the delayed write to storage devices. This is typically referred to as a write-back operation. FIG. 4d shows the writeback as done by the node with the working instance of the data blocks..."* The preceding text excerpt clearly indicates that the data blocks/write in data is written to storage devices after a delay/a predetermined time from the reception time spent in Node Y/temporary storage.) (Figure 4c; Figure 4d; Page 4, Paragraph 69).

Brown fails to disclose the storage device system further has a table in which the reception times corresponding to each write-in data which is waiting in the temporary storage device are arranged in the order from an older one.

Azagury discloses the storage device system further has a table in which the reception times corresponding to each write-in data which is waiting in the temporary storage device are arranged in the order from an older one (i.e. *" This information includes a timestamp provided by the host and a pointer to the modified data, which are in the control unit's cache. After the information is queued and normal write processing completes, the control unit signals the*

Art Unit: 2165

successful completion of the write to the host. The zSeries Parallel Sysplex has a timer facility, which enables a single clock to be shared among multiple zSeries hosts comprising the sysplex. Thus, by associating a sysplex timestamp with each write, it is possible to reconstruct the order in which the write requests were executed." The preceding text excerpt clearly indicates that an order/table is kept that arranged the timestamps/reception times of the write requests/write-in data in the order in which the write requests were executed/in the order from an older one.) (Page 6, Paragraph 4).

It would have been obvious in view of the prior art of record to modify the teachings of Brown with the teachings of Azagury to include discloses the storage device system further has a table in which the reception times corresponding to each write-in data which is waiting in the temporary storage device are arranged in the order from an older one with the motivation to enable the subsystem to support disaster recovery by ensuring that all data written to a primary control unit are also written to a secondary control unit (Azagury, Page 1, Paragraph 1).

As per Claim 5, Brown discloses a storage device system in a computer system having a plurality of superior devices and a plurality of storage device-systems for receiving write-in data from at least one of the superior devices, comprising: a physical storage device, responsive to a logical volume which is controlled so that identical data is saved across the plurality of storage device systems, for storing data on the logical volume (i.e. *"Volume managers are primarily used to organize storage devices into logical volumes, which may span multiple storage devices...Briefly, the present invention includes at least a first and a second server of a cluster of servers being equipped with complementary software RAID drivers and distributed lock managers that enable the first server to delegate to the second server, writing of a version of a unit of coherent data into a number of storage devices coupled to the server cluster."* The preceding text excerpt clearly indicates that a cluster of servers/a plurality of superior devices and a

Art Unit: 2165

second/delegated server and number of storage devices/a plurality of storage device systems exist that include a storage device/physical storage device which is responsive to a logical volume which is controlled with complementary software RAID drivers/controlled so that identical data is saved across the plurality of storage devices. Note that RAID arrays are used to save identical data across a plurality of storage devices.) (Page 1, Paragraphs 4, 15); a device for saving the time of reception on which write-in data was received from a superior device (i.e. *"in performing a delegated write, the delegated server may obtain at least a shared read lock on the unit of coherent data and validating a timestamp of the version of the unit of coherent data to be written."* The preceding text excerpt clearly indicates that delegated server obtains/validates/saves a timestamp/the time of reception on which write-in data was received from a first/ server/delegating server/superior device) (Pg 1, Paragraph 18); a communication interface device for transmitting write-in data addressed to the logical volume and a corresponding reception time to another storage device system and for receiving write-in data and corresponding reception time from the storage device system (i.e. *"...in performing a delegated write, the delegated server may obtain at least a shared read lock on the unit of coherent data and validating a timestamp of the version of the unit of coherent data to be written...the LastWriteTimeStamp property in the ManagedObject is the time at which the object was last written to the storage devices... FIGS. 4c and 4d show the delayed write to storage devices."* The preceding text excerpt clearly indicates that the delegated server validates the timestamp/reception time of the unit of data to be written. In order to validate the timestamp the delegating server must first receive it. Also note that because any of the cluster of servers can be designated as a delegating, or delegated server the server has the capacity to both transmit and receive the timestamp information. This property is demonstrated pertaining to the storage devices/storage device systems through the reference to the LastUpdateWriteStamp. The text excerpt also indicates, along with figures 4c and 4d, that write-in data may also be transferred in this way.) (Figures 4c, 4d; Pg 1, Paragraph 18; Page 3, Paragraph 46; Page 4, Paragraph 69); and a data consistency holding control device for effecting control with

reference to the table to write write-in data to the logical volume for which more than a predetermined time has passed from the reception time into the physical storage device in the order of the older reception time (i.e. *"FIGS. 4c and 4d show the delayed write to storage devices. This is typically referred to as a write-back operation. FIG. 4d shows the writeback as done by the node with the working instance of the data blocks..."*) The preceding text excerpt clearly indicates that the data blocks/write in data is written to storage devices after a delay/a predetermined time from the reception time spent in Node Y/temporary storage. Note that in order to determine the data to be written, the table would have to be referenced.) (Figure 4c; Figure 4d; Page 4, Paragraph 69).

Brown fails to disclose a table in which the reception times corresponding to each of write-in data from the superior device and write-in data from the other storage device system are arranged in order from an older one.

Azagury discloses a table in which the reception times corresponding to each of write-in data from the superior device and write-in data from the other storage device system are arranged in order from an older one (i.e. *" This information includes a timestamp provided by the host and a pointer to the modified data, which are in the control unit's cache. After the information is queued and normal write processing completes, the control unit signals the successful completion of the write to the host. The zSeries Parallel Sysplex has a timer facility, which enables a single clock to be shared among multiple zSeries hosts comprising the sysplex. Thus, by associating a sysplex timestamp with each write, it is possible to reconstruct the order in which the write requests were executed."*) The preceding text excerpt clearly indicates that an order/table is kept that arranged the timestamps/reception times of the write requests/write-in data in the order in which the write requests were executed/in the order from an older one.) (Page 6, Paragraph 4).

It would have been obvious in view of the prior art of record to modify the teachings of Brown with the teachings of Azagury to include a table in which the

Art Unit: 2165

reception times corresponding to each of write-in data from the superior device and write-in data from the other storage device system are arranged in order from an older one with the motivation to enable the subsystem to support disaster recovery by ensuring that all data written to a primary control unit are also written to a secondary control unit (Azagury, Page 1, Paragraph 1).

As per Claim 6, Brown discloses a storage device system in a computer system having a plurality of superior devices and a plurality of storage device-systems for receiving write-in data from at least one of the superior devices, comprising: a physical storage device, responsive to a logical volume which is controlled so that identical data is saved across the plurality of storage device systems, for storing data on the logical volume (i.e. *"Volume managers are primarily used to organize storage devices into logical volumes, which may span multiple storage devices...Briefly, the present invention includes at least a first and a second server of a cluster of servers being equipped with complementary software RAID drivers and distributed lock managers that enable the first server to delegate to the second server, writing of a version of a unit of coherent data into a number of storage devices coupled to the server cluster."* The preceding text excerpt clearly indicates that a cluster of servers/a plurality of superior devices and a second/delegated server and number of storage devices/a plurality of storage device systems exist that include a storage device/physical storage device which is responsive to a logical volume which is controlled with complementary software RAID drivers/controlled so that identical data is saved across the plurality of storage devices. Note that RAID arrays are used to save identical data across a plurality of storage devices.) (Page 1, Paragraphs 4, 15); a device for saving the time of reception on which write-in data was received from a superior device (i.e. *"in performing a delegated write, the delegated server may obtain at least a shared read lock on the unit of coherent data and validating a timestamp of the version of the unit of coherent data to be written."* The preceding text excerpt clearly

indicates that delegated server obtains/validates/saves a timestamp/the time of reception on which write-in data was received from a first/ server/delegating server/superior device) (Pg 1, Paragraph 18); a communication interface device for transmitting write-in data addressed to the logical volume and a corresponding reception time to another storage device system and for receiving write-in data and corresponding reception time from the storage device system (i.e. "...in performing a delegated write, the delegated server may obtain at least a shared read lock on the unit of coherent data and validating a timestamp of the version of the unit of coherent data to be written...the LastWriteTimeStamp property in the ManagedObject is the time at which the object was last written to the storage devices... FIGS. 4c and 4d show the delayed write to storage devices." The preceding text excerpt clearly indicates that the delegated server validates the timestamp/reception time of the unit of data to be written. In order to validate the timestamp the delegating server must first receive it. Also note that because any of the cluster of servers can be designated as a delegating, or delegated server the server has the capacity to both transmit and receive the timestamp information. This property is demonstrated pertaining to the storage devices/storage device systems through the reference to the LastUpdateWriteStamp. The text excerpt also indicates, along with figures 4c and 4d, that write-in data may also be transferred in this way.) (Figures 4c, 4d; Pg 1, Paragraph 18; Page 3, Paragraph 46; Page 4, Paragraph 69); a data consistency holding control device for effecting control with reference to the table to write write-in data to the logical volume for which more than a predetermined time has passed from the reception time into the physical storage device in the order of the older reception time (i.e. "FIGS. 4c and 4d show the delayed write to storage devices. This is typically referred to as a write-back operation. FIG. 4d shows the writeback as done by the node with the working instance of the data blocks..." The preceding text excerpt clearly indicates that the data blocks/write in data is written to storage devices after a delay/a predetermined time from the reception time spent in Node Y/temporary storage. Note that in order to determine the data to be written, the table would have to be referenced.) (Figure 4c; Figure 4d; Page 4, Paragraph 69) and the storage

Art Unit: 2165

device system further has a device for receiving a request for locking a partial region of the logical volume from the superior device and for locking the partial region, a device for transmitting the locking request which was received through the communication interface device to the other storage device system, a device for receiving the locking request through the communication interface device from the other storage device system and for locking a designated partial region and a device for rejecting a request for write-in of data from the superior device and the other storage device system to the partial region, except for a case in which it is a request from the superior device in which the partial region was locked (i.e. *"Once the writeback operation is started, the software RAID driver secures an exclusive write lock on the stripe."*) The preceding text excerpt clearly indicates that in order to secure the exclusive lock for the stripe/partial region, the second server/delegated server/storage system device must utilize the software RAID driver to request an exclusive lock from the first server/delegating server/superior device and then receive that request back in order to transmit the request to the logical volume, therefor locking the stripe/partial region. Note that an exclusive lock would deny access to the stripe to any system except the system which currently holds the lock.) (Page 4, Paragraph 71).

Brown fails to disclose a table in which the reception times corresponding to each of write-in data from the superior device and write-in data from the other storage device system are arranged in order from an older one.

Azagury discloses a table in which the reception times corresponding to each of write-in data from the superior device and write-in data from the other storage device system are arranged in order from an older one (i.e. *" This information includes a timestamp provided by the host and a pointer to the modified data, which are in the control unit's cache. After the information is queued and normal write processing completes, the control unit signals the successful completion of the write to the host. The zSeries Parallel Sysplex has a timer facility, which enables a*

single clock to be shared among multiple zSeries hosts comprising the sysplex. Thus, by associating a sysplex timestamp with each write, it is possible to reconstruct the order in which the write requests were executed." The preceding text excerpt clearly indicates that an order/table is kept that arranged the timestamps/reception times of the write requests/write-in data in the order in which the write requests were executed/in the order from an older one.) (Page 6, Paragraph 4).

It would have been obvious in view of the prior art of record to modify the teachings of Brown with the teachings of Azagury to include a table in which the reception times corresponding to each of write-in data from the superior device and write-in data from the other storage device system are arranged in order from an older one with the motivation to enable the subsystem to support disaster recovery by ensuring that all data written to a primary control unit are also written to a secondary control unit (Azagury, Page 1, Paragraph 1).

7. Claim 3 rejected under 35 U.S.C. 103(a) as being unpatentable over Brown (U.S. Pre Grant Publication Number 2003/0014598) in view of IBM ("Storage Management for IBM BS UDB: Snapshot Backup and Recovery With the IBM TotalStorage Enterprise Storage Server", 2002).

As per Claim 3, Brown discloses a storage device system in a computer system having a plurality of superior devices and a plurality of storage device systems for receiving write-in data from at least one of the superior devices, comprising: a physical storage device, responsive to a logical volume which is controlled so that identical data is saved across the plurality of storage device systems, for storing data on the logical volume (i.e. *"Volume managers are primarily used to organize storage devices into logical volumes, which may span multiple storage devices...Briefly, the present invention includes at least a first and a second server of a cluster of servers being equipped with complementary software RAID drivers and distributed lock managers that enable the first server to delegate to the second server, writing of a version*

of a unit of coherent data into a number of storage devices coupled to the server cluster." The preceding text excerpt clearly indicates that a cluster of servers/a plurality of superior devices and a second/delegated server and number of storage devices/a plurality of storage device systems exist that include a storage device/physical storage device which is responsive to a logical volume which is controlled with complementary software RAID drivers/controlled so that identical data is saved across the plurality of storage devices. Note that RAID arrays are used to save identical data across a plurality of storage devices.) (Page 1, Paragraphs 4, 15); a device for saving the time of reception on which write-in data was received from a superior device (i.e. *"...in performing a delegated write, the delegated server may obtain at least a shared read lock on the unit of coherent data and validating a timestamp of the version of the unit of coherent data to be written."* The preceding text excerpt clearly indicates that delegated server obtains/validates/saves a timestamp/the time of reception on which write-in data was received from a first/ server/delegating server/superior device) (Pg 1, Paragraph 18); a communication interface device for transmitting write-in data addressed to the logical volume and a corresponding reception time to another storage device system and for receiving write-in data and corresponding reception time from the storage device system (i.e. *"...in performing a delegated write, the delegated server may obtain at least a shared read lock on the unit of coherent data and validating a timestamp of the version of the unit of coherent data to be written...the LastWriteTimeStamp property in the ManagedObject is the time at which the object was last written to the storage devices... FIGS. 4c and 4d show the delayed write to storage devices."* The preceding text excerpt clearly indicates that the delegated server validates the timestamp/reception time of the unit of data to be written. In order to validate the timestamp the delegating server must first receive it. Also note that because any of the cluster of servers can be designated as a delegating, or delegated server the server has the capacity to both transmit and receive the timestamp information. This property is demonstrated pertaining to the storage devices/storage device systems through the reference to the LastUpdateWriteStamp. The text excerpt also indicates, along with figures 4c and 4d, that write-in data may also be transferred in this way.) (Figures 4c, 4d; Pg 1,

Paragraph 18; Page 3, Paragraph 46; Page 4, Paragraph 69); and a data consistency holding control device for effecting control to write write-in data which was received from the superior device and write-in data which was received through the communication interface device into the physical storage device after such write-in data has been made to wait in a temporary storage device for more than predetermined time from the reception time corresponding to the write-in data to the logical volume (i.e. *"FIGS. 4c and 4d show the delayed write to storage devices. This is typically referred to as a write-back operation. FIG. 4d shows the writeback as done by the node with the working instance of the data blocks..."*) The preceding text excerpt clearly indicates that the data blocks/write in data is written to storage devices after a delay/a predetermined time from the reception time spent in Node Y/temporary storage.) (Figure 4c; Figure 4d; Page 4, Paragraph 69), so that, when write-in data which was received from the superior device and write-in data which was received through the communication interface device are written in an overlapped manner into the same storage location of the physical storage device (note that such overlapped writing is also a feature of RAID), they are written in the order of the reception time thereof (i.e. *"...timestamp aging, as well as other techniques may be employed."*) The preceding text excerpt clearly indicates that timestamp aging, which assigns priority on a first come first serve basis, is used (e.g. they are written in the order of reception.) (Page 4, Paragraph 70), and a device for judging whether or not new write-in data is written in an overlapped manner into the same storage location as the other write-in data with reference to the bit map table (i.e. *"FIGS. 4c and 4d show the delayed write to storage devices. This is typically referred to as a write-back operation. FIG. 4d shows the writeback as done by the node with the working instance of the data blocks..."*) The preceding text excerpt clearly indicates that the data blocks/write in data is written to storage devices after a delay/a predetermined time from the reception

time spent in Node Y/temporary storage. Note that in order to determine the data to be written, the bitmap table would have to be referenced.) (Figure 4c; Figure 4d; Page 4, Paragraph 69).

Brown fails to disclose the storage device system further having a bit map table for setting a bit value to indicate whether or not each block of the write-in data exists in the temporary storage device.

IBM discloses the storage device system further having a bit map table for setting a bit value to indicate whether or not each block of the write-in data exists in the temporary storage device (i.e. *"...the ESS creates another copy of the data by building a bitmap that records changed data. When the bitmap is complete (typically in seconds), the copy is logically complete."* The preceding text excerpt clearly indicates that a bitmap/bitmap table is used to record changed data/indicate whether or not each block of the write-in data exists in the temporary storage device.) (Page 7, Paragraph 5).

It would have been obvious in view of the prior art of record to modify the teachings of Brown with the teachings of IBM to include the storage device system further having a bit map table for setting a bit value to indicate whether or not each block of the write-in data exists in the temporary storage device with the motivation to provide a synchronous remote mirror to protect from disasters (IBM, Page 4, Executive Summary).

Points Of Contact

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Michael J. Hicks whose telephone number is (571) 272-2670. The examiner can normally be reached on Monday - Friday 8:30a - 5:00p.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Jeffrey Gaffin can be reached on (571) 272-4146. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Michael J Hicks
Technology Center 2100
Art Unit 2165
(571) 272-2670



**SAM RIMELL
PRIMARY EXAMINER**